Compensating Demand Response Participants Via Their Shapley Values

Gearóid O'Brien, Abbas El Gamal and Ram Rajagopal

arXiv:1403.6713v1 [cs.GT] 26 Mar 2014

***Abstract*—Designing fair compensation mechanisms for demand response (DR) is challenging. This paper models the problem in a game theoretic setting and designs a payment distribution mechanism based on the Shapley Value. As exact computation of the Shapley Value is in general intractable, we propose estimating it using a reinforcement learning algorithm that approximates optimal stratified sampling. We apply this algorithm to two DR programs that utilize the Shapley Value for payments and quantify the accuracy of the resulting estimates.**

I. Introduction

Demand response, or adjusting the aggregate load profile as a means to help balance supply and demand in electricity grids, is becoming an important approach to improve grid reliability. Utilities such as PG&E [1] allow for third party operators to administer DR programs through schemes such as *Aggregator Managed Portfolios*. These operators (or 'aggregators') are responsible for most of the program details, including marketing, enrollment, and payments to participants.

Designing, implementing, and operating large scale DR programs is a non-trivial task, however. Ensuring that participants remain enrolled in the scheme — and that it is also appealing to new participants — relies in part on a fair and attractive compensation mechanism. The typical mechanism design problem focuses on distributing the total revenue accrued by an aggregator to all participants, in proportion to their contribution in addition to achieving fairness and enrollment incentives.

Design of demand response schemes has been extensively investigated in the literature [2]. Several papers have investigated price based mechanisms for various types of response capabilities [3], [4], [5], [6], [7], [8]. Mechanisms based on cooperative games have been investigated more recently [9], [10], but attention has focused on specific formats where consumers make choices under parameterized utility functions. Instead in this paper, we propose a simple payment scheme based on a traditional cooperative game solution concept: the Shapley Value.

Understanding the advantage of using the Shapley Value when compared to other conventional distribution methods is crucial to assessing the contribution of this paper. In the proposed mechanism each consumer that agrees to participate in a DR program receives a payment for cooperating and forming a coalition. The *value* of a coalition is a function measuring how close that coalition comes to achieving the goal of the DR program. A simple such function encapsulates the *penalty* (in dollars) imposed by the utility on the aggregator for failing to meet an agreed upon commitment for the chosen coalition. The aggregator needs to decide then how to distribute the penalty fairly among the participants. Distributing penalties or payments in DR schemes is integral to their success, primarily in situations where participants are free to choose from a number of schemes. The *Shapley Value* solution concept provides a fair and unique method for distributing the total penalty when the penalty function satisfies some conditions. Although the use of the word 'fair' may seem vague in this context, it is a precisely defined term satisfying the following four concepts, which reasonable distribution schemes should satisfy:

Efficiency: The entire payment or penalty is divided among the participants (no excess remains).
Symmetry: Two participants that contribute equally are rewarded equally.
Null Player: Participants that do not contribute receive no payoff.
Linearity: The total payoff rewarded for contributing to two programs is the sum of the payoff that would have been awarded for contributing to each of the two programs individually.

Surprisingly, the Shapley Value can be proven to be the *only* payment distribution method that satisfies the four axioms above, with the added benefit that the solution is unique. The issue of appropriately allocating penalties to participants is an important problem in demand response that has not been the focus of much research. The Shapley Value is an attractive solution to the problem.

The Shapley Value has been previously used in studies

G. O'Brien and A. El Gamal are with the Department of Electrical Engineering, Stanford University (e-mail: gobrien@stanford.edu,abbas@ee.stanford.edu). G. O'Brien is supported by a Stanford Graduate Fellowship and El Gamal is partially supported by the TomKat Center.

R. Rajagopal is with the Department of Civil and Environmental Engineering, Stanford University (e-mail: ramr@stanford.edu). Rajagopal is supported by a Powell Foundation Fellowship and the TomKat Center.

on electrical energy generation and transmission. In [11], the Shapley Value is used to allocate transmission service costs among network users in energy markets. In [12] the aggregation of wind power producers is studied using coalition game theory and show that the resulting game is not convex so the Shapley Value may not be appropriate.

The most challenging aspect in utilizing the Shapley Value is its computational intractability. For a DR program with n participants, the value function must be evaluated $n\,2^n$ times. A modest DR program with $n = 500$ requires 1.5×10^{153} function evaluations. Approximation approaches have been proposed to mitigate this problem. They rely on simple schemes to selectively perform function evaluations. Shapley proposed a Monte-Carlo random sampling technique [13], extended in [14] and [15] to achieve desired accuracy levels in polynomial time. Such mechanisms neither exploit relevant properties of the value function nor enforce important constraints such as budget balance.

This paper, which is an extended and more complete version of [16], proposes a Shapley Value based distribution of Demand Response payments (or penalties) and a new algorithm for estimation that is significantly faster and more accurate than prior approaches. The rest of this paper is organized as follows. Section II provides a brief introduction to the Shapley Value. Section III describes a method of estimating the Shapley Value using a pseudo-random sampling technique that significantly reduces the variance of the estimate when compared to random sampling. In Section IV, we analyze two simple demand response programs and utilize the Shapley Value as a means of compensating the participants in the scheme.

II. Demand Response and the Shapley Value

Consider a set $\mathcal{X} = \{1, 2, \ldots, n\}$ of n participants in a DR scheme. A "participant" could be a user or a user-load if such granularity is available. For the set $\mathcal{S} \subseteq \mathcal{X}$, define the value function (or characteristic function) $v(\mathcal{S})$ as the total penalty imposed on the participants in $\mathcal{S}$ if they do not achieve the DR goal. For example, if the goal of the DR program is to provide reserve by reducing load levels, the value function may be taken as

$$v(\mathcal{S}) = -\left[\sum_{i\in\mathcal{S}}\left(X_i - \tilde{X}_i\right)\right]_+, \tag{1}$$

where $X_i \in \mathbb{R}$ is the amount participant i agreed to reduce its load by and $\tilde{X}_i \in \mathbb{R}$ is the amount it actually reduced it by. As $[x]_+ = \max\{x, 0\}$, $v(\mathcal{S})$ is non-zero when the aggregate discrepancy is greater than 0. Additional choices of value functions are discussed in detail in Section IV. We assume the general function $v(\mathcal{S})$ is *submodular* (see [17] for a definition).

The operator wishes to distribute the total penalty $v(\mathcal{X})$ among the n participants in a fair manner, dependent on their relative contributions to the goal of the DR scheme. We denote the penalty assigned to participant i as ϕ_i. Hence the total penalty is

$$v(\mathcal{X}) = \sum_{i=1}^{n} \phi_i. \tag{2}$$

Shapley proposed a solution to the distribution of the total penalty that is both unique and fair when penalty functions are submodular [17] . For a given participant, it is the mean marginal contribution of that participant to all possible coalitions of the other participants. Defining $\mathcal{X}_{-i}$ to be the set of all participants after removing participant i, the marginal contribution of participant i to a coalition $\mathcal{S}$, $\mathcal{S} \subseteq \mathcal{X}_{-i}$, is

$$\rho_i(\mathcal{S}) = v(\mathcal{S} \cup \{i\}) - v(\mathcal{S}). \tag{3}$$

Furthermore, we define $\mathcal{R}$ to be one of the $n!$ permutations of the participants in $\mathcal{X}$, and $\mathcal{P}_i^{\mathcal{R}}$ to be the ordered set preceding i in $\mathcal{R}$. The Shapley Value is defined as

$$\phi_i(v) = \frac{1}{n!}\sum_{\mathcal{R}} \rho_i(\mathcal{P}_i^{\mathcal{R}}). \tag{4}$$

Clearly, direct calculation of the Shapley Value using equation (4) is intractable. For example, with a value function as seemingly elementary as that shown in (1), the function is of a form $\max\{x, 0\}$ and the Shapley value is computationally intensive to calculate. Hence, it is typically estimated using sampling techniques. The following section details a computationally efficient method for estimating the Shapley Value via sampling.

Remark. For brevity we suppress the v and denote the Shapley Value by ϕ_i. The value function assumed should be clear from the context.

A. Value Functions and Demand Response Schemes

In order to use the Shapley value as a distribution mechanism of a DR scheme, the scheme itself must be representable as a value function. This function, defined over subsets of the participants, returns the penalty that will be imposed on the operator of the DR scheme. Value functions are particularly suited to schemes where loads are controllable to some degree. The formulation of the value function is left to the designer of the DR scheme as this paper is neither involved with choosing the value function itself nor designing DR schemes.

III. Estimating the Shapley Value

Before describing our algorithm for estimating the Shapley Value, we need the following alternative formulation. Grouping the terms in equation (4) in which

the participants to the left of i are the same gives the alternative form for the Shapley Value

$$\phi_i = \sum_{\mathcal{S}\subseteq\mathcal{X}_{-i}} \frac{|\mathcal{S}|!(|\mathcal{X}|-|\mathcal{S}|-1)!}{n!}\rho_i(\mathcal{S}). \tag{5}$$

Further grouping by the number of terms in $\mathcal{S}$, defining $j=|\mathcal{S}|$, and recalling that $n=|\mathcal{X}|$ we obtain

$$\begin{aligned}\phi_i &= \sum_{j=0}^{n-1}\sum_{\mathcal{S}\subseteq\mathcal{X}_{-i}}\left(\frac{j!(n-j-1)!}{n!}\right)\rho_i(\mathcal{S})\\ &= \frac{1}{n}\sum_{j=0}^{n-1}\sum_{\mathcal{S}\subseteq\mathcal{X}_{-i}}\left(\frac{(n-1)!}{j!(n-1-j)!}\right)^{-1}\rho_i(\mathcal{S})\\ &= \frac{1}{n}\sum_{j=0}^{n-1}\sum_{\mathcal{S}\subseteq\mathcal{X}_{-i}}\binom{n-1}{j}^{-1}\rho_i(\mathcal{S}).\end{aligned}$$

The inner sum can be considered as an expectation over a uniform probability mass function, hence we can write

$$\phi_i = \frac{1}{n}\sum_{j=0}^{n-1}\mathbb{E}[\rho_i(\mathcal{S})]. \tag{6}$$

This form of the Shapley Value suggests an estimation approach based on stratified sampling [18]. For participant $i\in\{1,2,\dots,n\}$, let stratum j be the set of marginal contributions of that participant to every subset $\mathcal{S}\subseteq\mathcal{X}_{-i}$ of size $|\mathcal{S}|=j$. We randomly and independently draw N_j^i samples $\rho_{1,j}^i,\dots,\rho_{N_j^i,j}^i$ from each stratum j. Define the sample mean for participant i as the random variable

$$\begin{aligned}T(\rho_{k,j}^i) &= \frac{1}{n}\sum_{j=0}^{n-1}\frac{1}{N_j^i}\sum_{k=0}^{N_j^i}\rho_{k,j}^i\\ &= \frac{1}{n}\sum_{j=0}^{n-1}\overline{\rho}_j^i,\end{aligned} \tag{7}$$

where the random variable $\overline{\rho}_j^i$ is the sample mean of the data drawn from stratum j. This sample mean is a linear unbiased estimate of ϕ_i and would be a reasonable estimate of ϕ_i except for the fact that the sum of the estimates may not be equal to the total budget $v(\mathcal{X})$, which would violate the efficiency axiom of the Shapley Value. We therefore use the sample averages as basis for computing the maximum likelihood (ML) estimates of the Shapley Values with the budget constraint as follows.

Assume that the number of samples from each stratum is sufficiently large so that we can use the central limit theorem to approximate the distribution of $\overline{\rho}_j^i$ by a Gaussian with mean $\mu_{j,i}=\mathbb{E}[\rho_i(\mathcal{S})]$ and variance $\sigma_{j,i}^2$. By independence of the sample averages, it follows that the variance of $T(\rho_{k,j}^i)$,

$$\sigma_i^2 = \frac{1}{n^2}\sum_{j=0}^{n-1}\sigma_{j,i}^2,$$

and $T(\rho_{k,j}^i)\sim\mathcal{N}(\phi_i,\sigma_i^2)$.

By independence of the sample averages $f(T(\rho_{k,j}^i)|\phi_i)$, $i\in\mathcal{X}$, the likelihood function can be written as

$$f(T(\rho_{k,j}^1),\dots,T(\rho_{k,j}^n)|\phi_1\dots\phi_n) = \prod_{i=1}^{n} f\left(T(\rho_{k,j}^i)|\phi_i\right). \tag{8}$$

Since the sample averages $T(\rho_{k,j}^i)$ are Gaussian, we consider the log likelihood function

$$\sum_{i=1}^{n}\log\left(f\left(T(\rho_{k,j}^i)|\phi_i\right)\right) = \zeta - \sum_{i=1}^{n}\frac{\left(T(\rho_{k,j}^i)-\phi_i\right)^2}{2\sigma_i^2}, \tag{9}$$

where ζ is not a function of ϕ_i. To obtain the ML estimates of the Shapley Values we then need to solve the optimization problem with respect to ϕ_i:

$$\begin{aligned}&\underset{\phi_i}{\text{maximize}}\ \ \zeta - \sum_{i=1}^{n}\frac{\left(T(\rho_{k,j}^i)-\phi_i\right)^2}{2\sigma_i^2}\\ &\text{subject to}\ \ \sum_{i=1}^{n}\phi_i = v(\mathcal{X}).\end{aligned} \tag{10}$$

To solve this problem, we form the Lagrangian

$$\mathcal{L} = \zeta - \sum_{i=1}^{n}\frac{\left(T(\rho_{k,j}^i)-\phi_i\right)^2}{2\sigma_i^2} + \lambda\left(v(\mathcal{X}) - \sum_{i=1}^{n}\phi_i\right).$$

This is a convex optimization problem and has a simple analytical solution.

Theorem 3.1: The ML estimates of ϕ_i are given by

$$\hat{\phi}_i = T(\rho_{k,j}^i) - \frac{\sigma_i^2}{\sum_{m=0}^{n}\sigma_m^2}\big(\hat{v}(\mathcal{X}) - v(\mathcal{X})\big), \tag{11}$$

where $\hat{v}(\mathcal{X}) = \sum_{i=1}^{n}T(\rho_{k,j}^i)$.

Note that all properties of the Shapley Value (efficiency, symmetry, null player, linearity [19]) hold in expectation in equation (11), with the added benefit that the budget is *always* balanced, as the constraint in the optimization problem (10) is satisfied.

A. Sample allocation

We now turn our attention to the question of how many samples we should select from each stratum. Suppose we have a total budget of N samples per participant, i.e., $\sum_{j=0}^{n-1}N_j^i = N$ for every $i\in\mathcal{X}$. How do we divide them among the strata? One reasonable approach

would be to allocate the samples for each participant i to minimize the variance of the sample mean $T(\rho^i_{k,j})$ subject to $\sum_{j=0}^{n-1} N^i_j = N$. The following shows that the optimal sample allocation is the Neyman allocation [18] for equal weighting.

Lemma 3.2: The minimum variance of $T(\rho^i_{k,j})$ subject to $\sum_{j=0}^{n-1} N^i_j = N$ is

$$\sigma^2_{i,\mathrm{SD}} = \frac{1}{N}\mathrm{mean}(\sigma_{j,i})^2, \tag{12}$$

where $\sigma_{j,i}$ is the standard deviation of the population in stratum j for participant i. The value of $\mathrm{mean}(\sigma_{j,i})$ is calculated by averaging over the n values of $\sigma_{j,i}$ for participant i.

The values of N^i_j that achieve this minimum are

$$N^i_j = \frac{\sigma_{j,i}}{\sum_{m=0}^{n-1}\sigma_{m,i}},\ j \in \{0, 1, \ldots, n-1\}.$$

The proof of this lemma is given in the Appendix A.

It is interesting to compare the achievable variance of the sample means using the above optimal stratified sampling to the more commonly used uniform sampling. With uniform sampling, we draw N samples independently at random from the set of marginal contributions of participant i without taking strata into consideration. The variance of the sample average for this approach is

$$\sigma^2_{i,\mathrm{RS}} = \frac{1}{N}\left[\mathrm{mean}(\sigma^2_{j,i}) + \mathrm{var}(\mu_{j,i})\right], \tag{13}$$

where $\mu_{j,i}$ is the mean value of the population in stratum j for participant i. The value $\mathrm{var}(\mu_{j,i})$ is calculated as the variance of the n values of $\mu_{j,i}$ for participant i. The proof of this fact is given in Appendix B.

Sampling according to the Neyman Allocation (lemma 3.2) requires prior knowledge of the standard deviation of each stratum for each participant, which is not realistic. A more practical approach would be to sample equally from each of the n stratum, i.e., $N^i_j = N/n$. With this allocation, the variance of the sample average is

$$\sigma^2_{i,\mathrm{ES}} = \frac{1}{N}\mathrm{mean}(\sigma^2_{j,i}). \tag{14}$$

The proof of this fact follows readily from the proof of (13) and is omitted. Comparing the variances for these three allocation strategies, we can clearly see that

$$\sigma^2_{i,\mathrm{SD}} \leq \sigma^2_{i,\mathrm{ES}} \leq \sigma^2_{i,\mathrm{RS}}. \tag{15}$$

Hence, it is always better to sample in proportion to standard deviations. In the following section we describe a reinforcement learning algorithm for estimating these standard deviations during sampling.

B. Approximating Optimum Stratified Sampling

Implementing an approximation to SD sampling is a typical reinforcement learning problem in which the algorithm seeks to exploit the information it has about the standard deviations of the strata to sample correctly, but must at the same time explore in order to accurately calculate these very standard deviations. In our setting, the goal is to sample a specific (but unknown) number of times from each stratum. This differs from the usual reinforcement learning problems where the goal is to converge on a single optimum action that maximizes the total reward. This contrast means that some techniques (such as *ϵ-greedy*, *Pursuit* and *Reinforcement Comparison*) are not suitable, and other approaches must be altered to make them suitable for the problem at hand; see [20] for information on reinforcement learning. By comparison, stochastic methods [20] which assign a probability to each action in accordance with the expected reward (or standard deviation in this case) are quite suitable to our setting.

Our proposed algorithm 1 explicitly "explores" the problem space initially before gradually moving to an "exploit" phase in which it uses the results of the exploration to improve the sampling allocations. For participant i, the probability of sampling from stratum j at sample $t \leq N$ is

$$\pi_{j,i}(t) = \epsilon(t)\frac{1}{n} + (1-\epsilon(t))\frac{\hat{\sigma}_{j,i}}{\sum_{m=0}^{n-1}\hat{\sigma}_{m,i}}, \tag{16}$$

where $\hat{\sigma}_{j,i}$ is the current estimate of the standard deviation of stratum j. The choice of $\epsilon(t)$ is left to the user, but should be a decreasing function of t with $\epsilon(0) = 1$. We implemented a number of such functions (including the stepped function described in [21]) and found the most accurate to be the double sigmoid function

$$\epsilon(t) = \kappa - \frac{1}{1 + e^{-\frac{t-\gamma N}{\beta N}}}, \tag{17}$$

where κ is chosen to ensure $\epsilon(0) = 1$. Increasing γ in the above equation reduces the percentage of samples used for exploration, and increasing β increases the transition time from exploration to exploitation.

At each step t, Algorithm 1 chooses stratum j with probability $\pi_{j,i}(t)$ for participant i. The probabilities are then updated for the next iteration. The vector of standard deviations is updated in each step using a numerically stable algorithm from [22]. The algorithm returns the sample mean $T(\rho^i_{k,j})$ for participant i as well as the variance of that statistic, σ^2_i. Once this has been calculated for all n participants, the MLE can be computed using equation (11) to ensure that the budget is balanced.

If $\sigma^2_{i,\mathrm{SD}} << \sigma^2_{i,\mathrm{ES}}$, then implementing Algorithm 1 will significantly reduce the variance of the sample

Algorithm 1 Approximating SD sampling

procedure STANDARDDEVIATIONSAMPLING(N, i)
$t \leftarrow 1$
$\hat{\mu}_{j,i} \leftarrow \underline{0}$ ▷ Estimate of $\mu \in \mathbb{R}^n$
$\hat{\sigma}_{j,i} \leftarrow \underline{0}$ ▷ Estimate of $\sigma \in \mathbb{R}^n$
$c \leftarrow \underline{0}$ ▷ Vector where c_j is the number of samples taken from stratum j, $c \in \mathbb{R}^n$
$m2 \leftarrow \underline{0}$ ▷ Vector of the sum of squared differences from the current mean of stratum j, $m2 \in \mathbb{R}^n$
while $t \leq N$ **do**
$\pi_{j,i}(t) \leftarrow \epsilon(t)\frac{1}{n} + (1-\epsilon(t))\frac{\hat{\sigma}_j}{\sum_{m=0}^{n-1}\sigma_{m,i}}$
Choose stratum j at random, weighted by $\pi_{j,i}(t)$.
Choose a random coalition, $\mathcal{S} \subseteq \mathcal{X}_{-i}$ where $|\mathcal{S}| = j$.
$x \leftarrow \rho_i(\mathcal{S})$ ▷ x is a sample from stratum j
$c_j \leftarrow c_j + 1$ ▷ Update the count for stratum j
$\Delta \leftarrow x - \hat{\mu}_{j,i}$
$\hat{\mu}_{j,i} \leftarrow \hat{\mu}_{j,i} + \frac{\Delta}{c_j}$ ▷ Online update for estimate of $\mu_{j,i}$
$m2_j \leftarrow m2_j + \Delta(x - \hat{\mu}_{j,i})$
$\hat{\sigma}_{j,i} \leftarrow \sqrt{\frac{m2_j}{c_j - 1}}$ ▷ Online update for estimate of $\sigma_{j,i}$
end while
$T(\rho^i_{j,k}) \leftarrow \text{mean}(\hat{\mu_{j,i}})$
$\sigma_i^2 = \frac{1}{n^2}\sum_{j=0}^{n-1}\hat{\sigma}^2_{j,i}$
return $T(\rho^i_{j,k})$, σ_i^2
end procedure

mean. If however, $\sigma^2_{i,\text{SD}} \approx \sigma^2_{i,\text{ES}}$, then the benefit of the algorithm may well be outweighed by the complexity involved in the implementation and time involved in its execution. Comparing $\sigma^2_{i,\text{SD}}$ to $\sigma^2_{i,\text{ES}}$, we have

$$\frac{\sigma^2_{i,\text{SD}}}{\sigma^2_{i,\text{ES}}} = 1 + \frac{\text{var}(\sigma^2_{j,i})}{\text{mean}(\sigma_{j,i})^2}. \tag{18}$$

Hence, if $\text{var}(\sigma^2_{j,i})/\text{mean}(\sigma_{j,i})^2 \approx 0$, sampling equally from each strata would be preferable.

IV. DEMAND RESPONSE PROGRAMS

We demonstrate the use of the Shapley Value to distribute compensation among the participants in two illustrative DR programs. In practice, more complex value functions could be used achieving the similar gains in performance.

A. DR providing Reserve

In this DR example, each participant in the program agrees when requested to reduce its load by a predefined amount $X_i \in \mathbb{R}_+$ (loads can be reduced for example by dimming lights or controlling HVAC in a building). The operator then offers a quantity M of "spinning reserve" to a utility where $\sum_{i=1}^n X_i \geq M, M \in \mathbb{R}_+$. This ensures that the program has a leeway of $\Delta M = \sum_{i=1}^n X_i - M$. When a demand response event is requested, each participant responds appropriately. There may be a discrepancy between a participant's actual reduction in consumption, $\tilde{X}_i$, and the promised reduction amount X_i, which we denote $\Delta X_i = X_i - \tilde{X}_i$. The value ΔX_i can be thought of as participant i's contribution to the penalty which will be imposed if the total discrepancy across all n participants exceeds ΔM. We then take the value function for this DR program to be

$$v(\mathcal{S}) = -q\left[\sum_{i\in\mathcal{S}} \Delta X_i - \Delta M\right]_+, \tag{19}$$

where $\mathcal{S} \subseteq \mathcal{X}$ and $[x]_+ = \max\{x, 0\}$ (i.e., the value function is non-zero only when the DR program is unable to meet the agreed upon reduction of M) and $q > 0$ is a constant that converts energy to a penalty levied by the utility on the DR operator, which without loss of generality can be set to 1. Note that the value function in (19) is a form of a budget-additive function. It is submodular and as such is compatible with the Shapley Value.

A value function such as that in (19) models well the Aggregator Managed Portfolio (AMP) Demand Response programs offered through PG&E and operated by various third parties. Non-compliance penalties are imposed on the aggregators: "*The aggregators are penalized if they fail to deliver their committed load reductions. The penalties vary based on the shortfall, with larger penalties for larger shortfalls. Aggregators determine compensation and/or penalties for their participating customers.*" Using the Shapley Value as a

means of determining compensation is of great relevance to such aggregators.

To compare the performance of the sampling techniques we discussed in Section III, we consider a small set of $n = 20$ participants so that we can compute the exact Shapley Value (ground truth).

Figure 1 plots the sample mean and standard deviation for each stratum when using the value function (19) for a representative participant i. As can be seen, our stratified sampling algorithm which approximates sampling in proportion to the standard deviations shows significant improvements over both uniform and random sampling because strata 0 to 10 have zero mean (and standard deviation) and as such do not contribute to the Shapley Value and the samples taken from these strata in the uniform and random sampling methods are wasted.

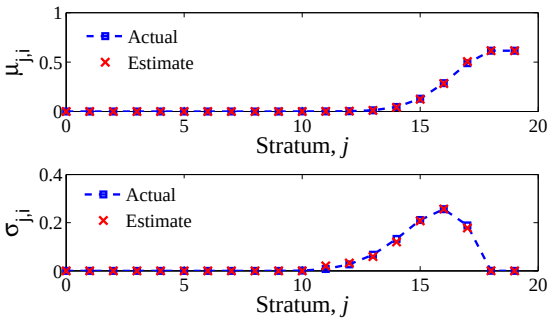

Fig. 1. Mean (top) and standard deviation (bottom) for each stratum. A red "x" indicates the final estimate from employing Algorithm 1 with $N = 5000$ samples.

Remark. As mentioned in Section III-B, we implemented the reinforcement learning algorithm using various $\epsilon(t)$ functions. Figure 2 plots the sample size against "regret," defined as the difference between the variance of the Shapley Value estimate for a given $\epsilon(t)$ and that of the estimate calculated using exact SD sampling. As can be seen, the sigmoid function (with $\gamma = 0.2$ and $\beta = 0.075$) which we use in all numerical results closely approximates ideal sampling.

Figure 3 shows the reduction in variance of the Shapley Value estimate as we change our sampling technique, indicating that estimating the standard deviations in this scenario could significantly reduce error in the Shapley Value. Implementing the proposed reinforcement learning algorithm reduces the empirical variance to a level approaching that achievable when the standard deviations are known in advance.

In Figure 4 we compare the estimated and actual Shapley Values for seven of the twenty participants for $0.7 \leq \Delta X_i \leq 0.8$. The estimates of the Shapley Values were calculated using $N = 5000$ samples. For this single instance, it can be seen that employing stratified sampling reduces the error significantly when compared to random sampling.

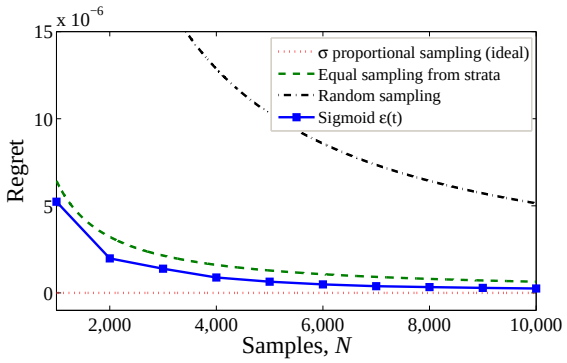

Fig. 2. Decrease in 'regret' as sample size increases for various $\epsilon(t)$ functions. The upper dashed line is the regret from uniform sampling and the lower dashed line at zero regret corresponds to ideal sampling.

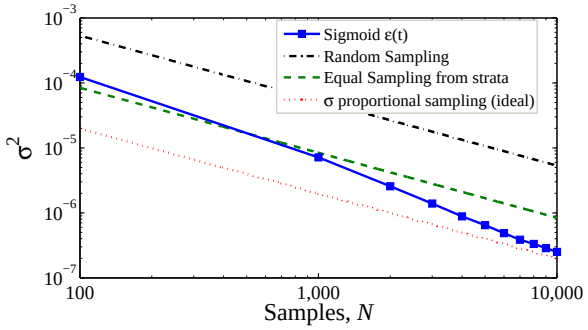

Fig. 3. Change in variance as sample size increases.

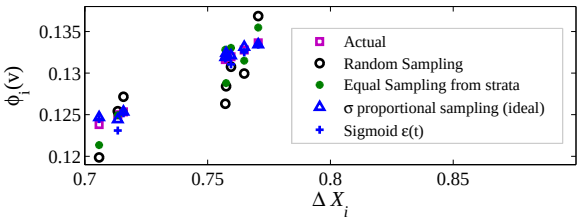

Fig. 4. Actual and estimates of the Shapley Values for seven of the twenty participants in the DR program. The estimates are calculated using $N = 5000$.

We determine the accuracy of the sampling methods using the mean squared prediction error

$$\text{MSPE} = \mathbb{E}\left[\left(\phi_i - \hat{\phi}_i\right)^2\right]. \tag{20}$$

For ease of comparison, we normalize the MSPE for each method by the MSPE for σ proportional sampling (the ideal method). Table I contains the comparison

TABLE I
THE NORMALIZED MSPE FOR THE VARIOUS SAMPLING METHODS.

Method	Normalized MSPE
Random Sampling	26.3084
Equal Sampling from strata	4.6521
σ proportional sampling (ideal)	1
Sigmoid $\epsilon(t)$	1.8050

results. It is clear that employing stratified sampling gives much better results than simple random sampling. The learning algorithm significantly outperforms uniform stratified sampling and approaches the accuracy of ideal stratified sampling.

B. Deferrable Load following

The second DR program we consider is deferrable load following, which is described in detail in [23]. In this program, $\boldsymbol{X}_i \in \mathbb{R}_+^T$ is a load profile for participant $i \in \mathcal{X}$ of T time steps. Given a set $\mathcal{S} \subseteq \mathcal{X}$, the operator wishes to *schedule* (i.e., delay in time) a number of loads of the participants in $\mathcal{S}$ such that the new aggregate load profile $\mathbf{s} \in \mathbb{R}_+^T$ approximates a predefined target load profile $\mathbf{y} \in \mathbb{R}_+^T$. This target profile is chosen in advance. Each load has a maximum allowable delay which may be 0 if the load is not schedulable. This problem reduces to a "knapsack packing" exercise, hence it is NP hard. To approximate a polynomial time solution, we use a greedy algorithm [23] that analyses each load and its set of possible delays in order to maximize at each step the functional

$$v(\mathcal{S}) = \frac{1}{T}\left(||\mathbf{y}||_2^2 - ||\mathbf{y} - \mathbf{s}||_2^2\right), \qquad (21)$$

where $\mathcal{S} \subseteq \mathcal{X}$, and $\mathbf{s}$ is the aggregate load profile produced by appropriately scheduling the loads of the participants in $\mathcal{S}$. The function $v(\mathcal{S})$ is maximized when $\mathbf{s} = \mathbf{y}$. The interested reader is referred to [23] for a detailed discussion of this scheme. In Figure 5 we see the unscheduled aggregate together with the target profile $\mathbf{y}$ on the left and the scheduled aggregate $\mathbf{s}$ on the right. The data used to generate these figures and results was taken from the *Plugwise* dataset, which contains plug level load data. The raw data is divided into 24 hour blocks. Each block is labeled with a user ID, timestamp, device description, and total energy consumption per hour for that device.

We assume that the operator of the DR program has realized revenue from operating the scheme and wishes to distribute this revenue fairly among the participants using the Shapley Value.

The nature of the load following algorithm means that it is not very illuminating to analyze small load sets. Hence in this case, we take a much larger load set containing $n = 500$ load profiles. For the load

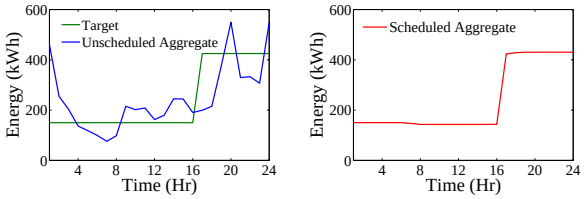

Fig. 5. Unscheduled aggregate on the left and scheduled aggregate on the right.

following scheme, each load profile has 24 data points corresponding to hourly smart-meter readings. Again, a typical load was isolated from the Plugwise dataset and a detailed analysis was performed on that load.

Figure 6 plots the sample means and standard deviations of the strata when using the value function (21). Calculating equation (18) for this instance results in $\sigma^2_{i,\mathrm{SD}} \approx \sigma^2_{i,\mathrm{ES}}$, and so the benefits of algorithm 1 will not be significant in this case.

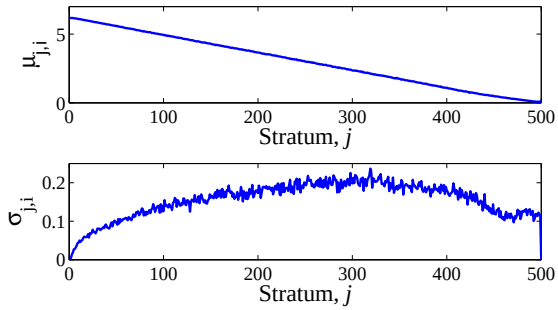

Fig. 6. Estimated mean (top) and estimated standard deviation (bottom) for each stratum. Each stratum was sampled 200 times to estimate the mean and standard deviation.

Using the estimates of the stratum means and standard deviations in figure 6 to calculate the variance of the Shapley Value estimated using the three sampling techniques produces plots in figure 7. As the plots for change in variance using both ideal sampling and uniform weighted sampling are quite similar, the added benefit of using algorithm 1 will indeed be outweighed by the complexity involved in its implementation.

V. CONCLUSION

This paper proposes the use of the Shapley Value to distribute the penalty among the participants in a DR program. As the Shapley Value is computationally intractable in general, we proposed a stratified sampling technique that reduces the number of samples needed to achieve a desired estimation accuracy while satisfying the budget balance constraint. We found that optimal stratified sampling requires prior knowledge of

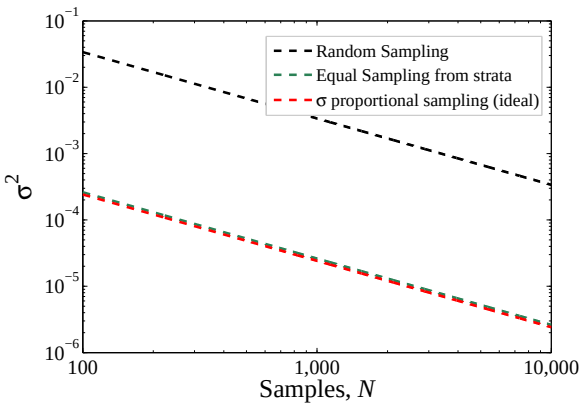

Fig. 7. Change in variance as sample size increases for the deferrable load scheduling program. Note how close the curves are for both ideal sampling and uniform weighted sampling.

the standard deviations of the strata, which may not be available. As such, we proposed a reinforcement learning heuristic which estimates the standard deviations and uses them to adjust the sample allocation among the strata. We demonstrated the use of the Shapley Value in DR programs numerically, describing one scenario (DR providing reserve) where the reinforcement learning algorithm can significantly reduce the variance of the estimate and another scenario (deferrable load following) where sampling equally from each stratum is very nearly as effective as implementing the algorithm.

It should be noted also that this method is agnostic to the specifics of the characteristic function under analysis, and can therefore be used to estimate the Shapley Value for any cooperative games, not only the DR programs analyzed in this paper. Also, this method ensures that the "budget balancing" constraint is met. To our knowledge, this constraint has not previously been considered in other research on estimating the Shapley Value using random sampling techniques. However, its importance is clear in DR programs such as those described here, where a given penalty needs to be distributed in its entirety among participants.

APPENDIX

A. Proof of Lemma 3.2

The variance of the sample mean of stratum j is

$$\mathrm{var}(\hat{\mu}_{j,i}) = \frac{\sigma_{j,i}^2}{N_{j,i}}.$$

The optimization problem is therefore:

$$\begin{aligned} \text{minimize} \quad & \mathrm{var}(T(\rho_{k,j}^i)) \\ \text{subject to} \quad & \sum_{j=0}^{n-1} N_{j,i} = N. \end{aligned}$$

Forming the Lagrangian

$$\mathcal{L} = \frac{1}{n^2}\sum_{j=0}^{n-1}\frac{\sigma_{j,i}^2}{N_{j,i}} + \lambda\left(\sum_{j=0}^{n-1} N_{j,i} - N\right).$$

Differentiating and setting equal to 0 results in

$$N_{j,i} = N\frac{\sigma_{j,i}}{\sum_{m=0}^{n-1}\sigma_{m,i}}.$$

Substituting back yields

$$\mathrm{var}(T(\rho_{k,j}^i)) = \frac{1}{N}\left(\frac{1}{n}\sum_{j=0}^{n-1}\sigma_{j,i}\right)\cdot\left(\frac{1}{n}\sum_{m=0}^{n-1}\sigma_{m,i}\right),$$
$$\mathrm{var}(\hat{\phi}_i) = \frac{1}{N}\left(\mathrm{mean}(\sigma_{j,i})\right)^2.$$

B. Proof of Equation (13)

A formulation for the Shapley Value is (see equation (4))

$$\phi_i = \frac{1}{n!}\sum_R \rho_i(P_i^R),$$

where R is an ordering of all players and P_i^R is the set of players which precede i in the order R. The statistic $T(\rho_{k,j}^i)$ is calculated by a random sampling of orderings R, i.e., if N random orderings R are analyzed, the statistic is

$$T(\rho_{k,j}^i) = \frac{1}{N}\sum_R \rho_i(P_i^R).$$

Assuming the marginal contributions are uncorrelated, the variance of the estimate is

$$\mathrm{var}\left(T(\rho_{k,j}^i)\right) = \frac{1}{N^2}\sum_R \mathrm{var}\left(\rho_i(P_i^R)\right).$$

We define the random variable $\mathcal{J}$ to be a discrete uniform distribution with support $\{0,1,2,\ldots,n-1$, and a sample from this distribution is therefore a random strata index, j. By the law of total variance and conditioning on $\mathcal{J}$, we then have

$$\begin{aligned}\mathrm{var}\left(\rho_i\left(P_i^R\right)\right) = {} & \mathbb{E}\left[\mathrm{var}\left(\rho_i\left(P_i^R\right)|\mathcal{J}\right)\right] \\ & + \mathrm{var}\left(\mathbb{E}\left[\rho_i\left(P_i^R\right)|\mathcal{J}\right]\right).\end{aligned}$$

We note that

$$\sigma_{j,i}^2 = \mathrm{var}\left(\rho_i\left(P_i^R\right)|\mathcal{J}\right) \text{ and } \mu_{j,i} = \mathbb{E}\left[\rho_i\left(P_i^R\right)|\mathcal{J}\right].$$

Therefore,

$$\mathrm{var}\left(\rho_i\left(P_i^R\right)\right) = \mathrm{mean}(\sigma_{j,i}^2) + \mathrm{var}\left(\mu_{j,i}\right).$$

Substituting back gives

$$\begin{aligned}\mathrm{var}\left(T(\rho_{k,j}^i)\right) &= \frac{1}{N^2}\sum_R\left[\mathrm{mean}(\sigma_{j,i}^2) + \mathrm{var}\left(\mu_{j,i}\right)\right] \\ &= \frac{1}{N}\left[\mathrm{mean}(\sigma_{j,i}^2) + \mathrm{var}\left(\mu_{j,i}\right)\right].\end{aligned}$$

REFERENCES

[1] PG&E, "Energy Management Programs," http://www.pge.com/en/mybusiness/save/energymanagement/index.page, [Online; accessed 25-July-2013].

[2] S. Borenstein *et al.*, "Dynamic pricing, advanced metering, and demand response in electricity markets," Center for the Study of Energy Markets., Tech. Rep. 3, Oct. 2002.

[3] N. Li *et al.*, "Optimal demand response based on utility maximization in power networks," in *Power and Energy Society General Meeting, 2011 IEEE*, 2011, pp. 1–8.

[4] N. Gatsis and G. Giannakis, "Cooperative multi-residence demand response scheduling," in *Information Sciences and Systems (CISS), 2011 45th Annual Conference on*, 2011, pp. 1–6.

[5] A. Conejo *et al.*, "Real-time demand response model," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 236–242, 2010.

[6] D. O'Neill *et al.*, "Residential demand response using reinforcement learning," in *1st IEEE Int. Conf. Smart Grid Communications*, 2010, pp. 409–414.

[7] Q. Dong *et al.*, "Distributed demand and response algorithm for optimizing social-welfare in smart grid," in *IEEE 26th Int. Parallel Distributed Processing Symposium (IPDPS)*, 2012, pp. 1228–1239.

[8] T. Kim and H. Poor, "Scheduling power consumption with price uncertainty," *IEEE Trans. Smart Grid*, vol. 2, no. 3, pp. 519–527, 2011.

[9] Q. Zhu *et al.*, "A game-theoretic framework for control of distributed renewable-based energy resources in smart grids," in *American Control Conference (ACC), 2012*, 2012, pp. 3623–3628.

[10] W. Saad *et al.*, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *IEEE Signal Processing Mag.*, vol. 29, no. 5, pp. 86–105, 2012.

[11] M. Junqueira *et al.*, "An Aumann-Shapley approach to allocate transmission service cost among network users in electricity markets," *IEEE Trans. Power Syst.*, vol. 22, no. 4, pp. 1532–1546, 2007.

[12] E. Baeyens *et al.*, "Wind energy aggregation: A coalitional game approach," in *50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*. IEEE, 2011, pp. 3000–3007.

[13] I. Mann and L. S. Shapley, "Values for large games, IV: Evaluating the Electoral College by Monte Carlo Techniques," The Rand Corporation, Tech. Rep., 1960.

[14] Y. Bachrach *et al.*, "Approximating power indices: theoretical and empirical analysis," *Autonomous Agents and Multi-Agent Systems*, vol. 20, no. 2, pp. 105–122, Mar 2010.

[15] J. Castro *et al.*, "Polynomial calculation of the shapley value based on sampling," *Computers & Operations Research*, vol. 36, no. 5, pp. 1726–1730, 2009.

[16] G. O'Brien *et al.*, "Efficient computation of shapley values for demand response programs," in *Proc. of the fourth Int. Conf. on Future energy systems*, ser. e-Energy '13. ACM, 2013, pp. 283–284.

[17] L. S. Shapley, "A Value for N-Person Games," *Contributions to the Theory of Games*, vol. 2, pp. 31–40, 1953.

[18] P. S. Rao, *Sampling Methodologies*. Chapman & Hall/CRC, 2000.

[19] G. Owen, *Game Theory*, 2nd ed. Academic Press, 1982.

[20] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge Univ Press, 1998, vol. 1, no. 1.

[21] P. Etoré and B. Jourdain, "Adaptive optimal allocation in stratified sampling methods," *Methodology and Computing in Applied Probability*, vol. 12, no. 3, pp. 335–360, 2010.

[22] D. E. Knuth, *The art of computer programming, volume 2 (3rd ed.): seminumerical algorithms*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1997.

[23] G. O'Brien and R. Rajagopal, "A Method for Automatically Scheduling Notified Deferrable Loads," in *American Control Conference (ACC)*, June 2013.

arXiv:1403.6713v1 [cs.GT] 26 Mar 2014

APPENDIX

A. Proof of Lemma 3.2

The variance of the sample mean of stratum j is

$$\mathrm{var}(\hat{\mu}_{j,i}) = \frac{\sigma^2_{j,i}}{N_{j,i}}.$$

The optimization problem is therefore:

$$\begin{aligned} &\text{minimize } \mathrm{var}(T(\rho^i_{k,j})) \\ &\text{subject to } \sum_{j=0}^{n-1} N_{j,i} = N. \end{aligned}$$

Forming the Lagrangian

$$\mathcal{L} = \frac{1}{n^2}\sum_{j=0}^{n-1}\frac{\sigma^2_{j,i}}{N_{j,i}} + \lambda\left(\sum_{j=0}^{n-1} N_{j,i} - N\right).$$

Differentiating and setting equal to 0 results in

$$N_{j,i} = N\frac{\sigma_{j,i}}{\sum_{m=0}^{n-1}\sigma_{m,i}}.$$

Substituting back yields

$$\begin{aligned} \mathrm{var}(T(\rho^i_{k,j})) &= \frac{1}{N}\left(\frac{1}{n}\sum_{j=0}^{n-1}\sigma_{j,i}\right)\cdot\left(\frac{1}{n}\sum_{m=0}^{n-1}\sigma_{m,i}\right), \\ \mathrm{var}(\hat{\phi}_i) &= \frac{1}{N}\left(\mathrm{mean}(\sigma_{j,i})\right)^2. \end{aligned}$$

B. Proof of Equation (12)

A formulation for the Shapley Value is (see equation (3))

$$\phi_i = \frac{1}{n!}\sum_R \rho_i(P^R_i),$$

where R is an ordering of all players and P^R_i is the set of players which precede i in the order R. The statistic $T(\rho^i_{k,j})$ is calculated by a random sampling of orderings R, i.e., if N random orderings R are analyzed, the statistic is

$$T(\rho^i_{k,j}) = \frac{1}{N}\sum_R \rho_i(P^R_i).$$

Assuming the marginal contributions are uncorrelated, the variance of the estimate is

$$\mathrm{var}\left(T(\rho^i_{k,j})\right) = \frac{1}{N^2}\sum_R \mathrm{var}\left(\rho_i(P^R_i)\right).$$

We define the random variable $\mathcal{J}$ to be a discrete uniform distribution with support $\{0, 1, 2, \ldots, n-1$, and a sample from this distribution is therefore a random strata index, j. By the law of total variance and conditioning on $\mathcal{J}$, we then have

$$\begin{aligned} \mathrm{var}\left(\rho_i\left(P^R_i\right)\right) = &\mathbb{E}\left[\mathrm{var}\left(\rho_i\left(P^R_i\right)|\mathcal{J}\right)\right] \\ &+ \mathrm{var}\left(\mathbb{E}\left[\rho_i\left(P^R_i\right)|\mathcal{J}\right]\right). \end{aligned}$$

We note that

$$\sigma^2_{j,i} = \mathrm{var}\left(\rho_i\left(P^R_i\right)|\mathcal{J}\right) \text{ and } \mu_{j,i} = \mathbb{E}\left[\rho_i\left(P^R_i\right)|\mathcal{J}\right].$$

Therefore,

$$\mathrm{var}\left(\rho_i\left(P^R_i\right)\right) = \mathrm{mean}(\sigma^2_{j,i}) + \mathrm{var}\left(\mu_{j,i}\right).$$

Substituting back gives

$$\begin{aligned} \mathrm{var}\left(T(\rho^i_{k,j})\right) &= \frac{1}{N^2}\sum_R\left[\mathrm{mean}(\sigma^2_{j,i}) + \mathrm{var}\left(\mu_{j,i}\right)\right] \\ &= \frac{1}{N}\left[\mathrm{mean}(\sigma^2_{j,i}) + \mathrm{var}\left(\mu_{j,i}\right)\right]. \end{aligned}$$

C. Proof of Equation (13)

For participant i, the variance of the sample mean of stratum j is

$$\mathrm{var}(\hat{\mu}_{j,i}) = \frac{\sigma^2_{j,i}}{N/n}.$$

As $T(\rho^i_{k,j})$ is the mean of $\hat{\mu_{j,i}}$ and assuming the strata are uncorrelated, we have

$$\mathrm{var}(T(\rho^i_{k,j})) = \frac{1}{n^2}\sum_{j=0}^{n-1}\mathrm{var}(\hat{\mu}_{j,i}) = \frac{1}{N}\mathrm{mean}(\sigma^2_{j,i}).$$